

WHAT IS CLAIMED IS:

1. A method of determining a motional state of a mobile robot, the method comprising:

retrieving pixel data for images taken at intervals from a camera that is coupled to the mobile robot;

comparing pixel data for a first image to pixel data for a second image to generate a measure of a difference between the two images, wherein comparing comprises:

filtering the first image pixel data with a gradient magnitude filter, where the gradient magnitude filter computes at least a spatial gradient;

comparing the gradient-magnitude filtered first image pixel data to a first threshold;

generating a binary map of the first image pixel data at least partly in response to the comparison of the gradient-magnitude filtered smoothed first image pixel data with the first threshold;

filtering the second image pixel data with the gradient magnitude filter;

comparing the gradient-magnitude filtered smoothed second image pixel data to a second threshold;

generating a binary map of the second image pixel data at least partly in response to the comparison of the gradient-magnitude filtered second image pixel data to the second threshold; and

comparing the binary map of the first image pixel data to the binary map of the second image pixel data to identify data for pixels that are different between the first image and the second image;

using the comparison of the pixel data to count the number of pixel data identified as changed;

comparing the count to a third predetermined threshold; and

determining the motional state of the mobile object at least partly in response to the count.

2. The method as defined in Claim 1, wherein the method is performed in real time.

3. The method as defined in Claim 1, wherein the gradient magnitude filter further comprises a low-pass filter.

4. The method as defined in Claim 3, wherein the low-pass filter corresponds to a Gaussian filter.

5. The method as defined in Claim 3, wherein filtering the first image pixel data further comprises:

filtering the first image pixel data in a first direction and in a second direction with a gradient filter that is characterized by the following convolution kernel, $[-1 \ -4 \ -5 \ 0 \ 5 \ 4 \ 1] / 32$, where results of filtering are separately maintained for the first direction and for the second direction; and

combining the separately maintained data by computing a magnitude of the data.

6. The method as defined in Claim 1, further comprising low-pass filtering image pixel data prior to computing a gradient.

7. The method as defined in Claim 6, wherein a low-pass filter is characterized by the following convolution kernel: $[1 \ 4 \ 6 \ 4 \ 1] / 16$; and a gradient filter is characterized by the following convolution kernel: $[-1 \ 0 \ 1] / 2$.

8. The method as defined in Claim 1, further comprising:
receiving an indication that the mobile robot is intended to be in motion;
determining that the motional state of the mobile robot is not in motion; and
providing an indication to a user that an undesired stoppage of the mobile robot has occurred.

9. The method as defined in Claim 1, further comprising:
receiving an indication that the mobile robot is intended to be in motion in a first direction;
determining that the motional state of the mobile robot is not in motion while receiving the indication that the mobile robot is intended to be in motion; and
changing a path of travel for the mobile robot at least partly in response to the determination.

10. The method as defined in Claim 1, further comprising subsampling the pixel data such that fewer than all the available pixels from the camera are used.

11. The method as defined in Claim 1, wherein the third predetermined threshold varies at least partly in response to a number of pixels used to compute the count.

12. The method as defined in Claim 1, wherein the first threshold and the second threshold are adaptive to filtering by the gradient magnitude filter of the first image pixel data and the second image pixel data, respectively.

13. The method as defined in Claim 12, wherein the first threshold is selected to be about half the maximum value for the gradient magnitude computed for a pixel by the gradient magnitude filter after filtering of the first image pixel data, and where the second threshold is selected to be about half the maximum value for the gradient magnitude computed for a pixel by the gradient magnitude filter after filtering of the second image pixel data.

14. The method as defined in Claim 1, wherein the comparison of the binary map of the first image pixel data to the binary map of the second image pixel data comprises a sloppy exclusive-or (XOR) computation.

15. The method as defined in Claim 14, wherein the sloppy XOR is computed by comparing a pixel from the binary map of the first image pixel data with a location of (x,y) to a plurality of pixels in the binary map of the second image pixel data, wherein the plurality of pixels includes a pixel with a location of (x,y).

16. The method as defined in Claim 1, wherein the motional state is determined only with observation of visual data.

17. The method as defined in Claim 1, where the mobile robot is autonomous, and where the method is performed entirely within the mobile robot.

18. The method as defined in Claim 1, further comprising:
sending pixel data from the mobile robot to a remote computer;
performing at least part of filtering operations in the remote computer; and
receiving an indication of the motional state of the mobile robot from the remote computer.

19. A method of determining a motional state of a mobile robot, the method comprising:

receiving pixel data for video images, where the video images are taken from a camera mounted to the mobile robot;

processing the pixel data for the video images to identify amounts of spatial gradient within a video image;

characterizing pixels of a video image into at least a first group and a second group, wherein the pixels of the first group correspond to a higher spatial gradient than the pixels of the second group; and

using the characterization of the pixels to compare a first video image to a second video image to detect the motional state of the mobile robot.

20. The method as defined in Claim 19, wherein characterizing further comprises:
arranging the pixels of a video image into a binary map such that a first logic state indicates association with the first group with the higher spatial gradient; and
a second logic state indicates association with the second group.

21. The method as defined in Claim 20, wherein using the characterization comprises:

comparing a binary value for one pixel from the binary map of the first image to a group of binary values for a group of pixels in the second image;

wherein the group of pixels in the second image includes a pixel in the same location as the one pixel from the binary map of the first image.

22. The method as defined in Claim 21, where comparing the binary value to the group of binary values further comprises computing a sloppy Exclusive-OR (“sloppy XOR”) operation between the binary value and the group of binary values.

23. The method as defined in Claim 19, further comprising:

inspecting pixel data to evaluate whether enough useful spatial gradients exist for robust detection of the motional state; and

inhibiting a motional state of “not moving” at least partly in response to a determination that the detection of the motional state is not likely to be reliable.

24. The method as defined in Claim 19, wherein the motional state is determined only by analysis of visual data observed by the camera.

25. The method as defined in Claim 19, where the mobile robot is autonomous, and where the method is performed entirely within the mobile robot.

26. The method as defined in Claim 19, further comprising:
sending pixel data from the mobile robot to a remote computer;
performing at least part of pixel characterizing operations in the remote computer; and
receiving an indication of the motional state of the mobile robot from the remote computer.

27. A method of controlling a behavior of a mobile robot based on a mismatch between an intended motional state and a perceived motional state in a mobile robot, the method comprising:

receiving an indication of the intended motional state, where the motional state is selected from the group including moving and not moving;

using visual data from a camera that is coupled to the mobile robot to perceive the motional state of the mobile robot, where the perceived motional state of the mobile robot is selected from the group including moving and not moving;

comparing the intended motional state to the perceived motional state to detect whether a mismatch exists between the intended motional state and the perceived motional state; and

changing the behavior of the mobile robot at least partly in response to a detected mismatch.

28. The method as defined in Claim 27, wherein the intended motional state is moving, where changing the behavior comprises changing a navigated path, further comprising:

instructing the mobile robot to travel in a direction approximately opposite to that previously requested for at least a distance sufficient to clear an obstruction;

instructing the mobile robot to yaw; and

instructing the mobile robot to resume traveling.

29. The method as defined in Claim 27, wherein the intended motional state is moving, and where changing the behavior comprises shutting off motors to conserve battery life.

30. The method as defined in Claim 27, wherein the mobile robot further comprises floor cleaning brushes, wherein changing the behavior further comprises shutting off cleaning brushes.

31. The method as defined in Claim 27, wherein the mobile robot further comprises a vacuum cleaner, wherein changing the behavior further comprises shutting off the vacuum cleaner.

32. The method as defined in Claim 27, wherein the intended motional state is moving, and where changing the behavior comprises setting an alert.

33. The method as defined in Claim 27, wherein the intended motional state is moving, and where changing the behavior comprises setting an alert.

34. The method as defined in Claim 27, wherein the mobile robot uses simultaneous localization and mapping (SLAM) techniques for navigation, where the intended motional state is not moving, wherein changing the behavior comprises:

resetting an estimate of pose (location and heading); and

disabling mapping functions for recognition of new landmarks until localization is achieved.

35. The method as defined in Claim 27, wherein the mobile robot uses simultaneous localization and mapping (SLAM) techniques for navigation, where the intended motional state is not moving, wherein changing the behavior comprises automatically changing the motional state to moving.

36. The method as defined in Claim 35, wherein changing the behavior further comprising resetting an estimate of pose, including location and heading.

37. The method as defined in Claim 27, wherein the mobile robot uses simultaneous localization and mapping (SLAM) techniques for navigation, where the intended motional state is not moving, wherein changing the behavior comprises disabling mapping functions for recognition of new landmarks until passage of a predetermined amount of time.

38. The method as defined in Claim 37, wherein changing the behavior further comprising resetting an estimate of pose, including location and heading.

39. The method as defined in Claim 27, wherein the mobile robot uses simultaneous localization and mapping (SLAM) techniques for navigation, where the

intended motional state is not moving, wherein changing the behavior comprises disabling mapping functions for recognition of new landmarks until a predetermined number of unmatched landmarks have been observed.

40. The method as defined in Claim 39, wherein changing the behavior further comprising resetting an estimate of pose, including location and heading.

41. The method as defined in Claim 27, wherein the intended motional state is moving, where changing the behavior comprises:

using the indication to determine that the mobile robot has been knocked over;
and

initiating procedures to restore the mobile robot to an upright position.

42. The method as defined in Claim 27, wherein the motional state is perceived using only the visual data from the camera.

43. The method as defined in Claim 27, where the mobile robot is autonomous, and where the method is performed entirely within the mobile robot.

44. The method as defined in Claim 27, further comprising:

transferring visual data from the mobile robot to a remote computer;

performing at least part of filtering operations in the remote computer; and

transferring an indication of the perceived motional state of the mobile robot from the remote computer.

45. A method of controlling the motion of a self-navigating mobile robot, the method comprising:

receiving an indication that the mobile robot is intended to be traveling in a forward direction;

determining from visual image data collected from a video camera coupled to the mobile robot that the mobile robot has ceased traveling in a forward direction;

discontinuing commands to propel the mobile robot in the forward direction;

commanding the mobile robot to travel in a reverse direction for at least a predetermined distance;

determining that the mobile robot has traveled in the reverse direction for at least about the predetermined distance;

discontinuing commands to propel the mobile robot in the reverse direction;

instructing the mobile robot to yaw by at least a first predetermined angle; and
commanding the mobile robot to resume forward motion.

46. The method as defined in Claim 45, wherein the predetermined distance is about 0.2 meters.

47. The method as defined in Claim 45, wherein the predetermined distance corresponds to at least an amount sufficient to permit the mobile robot to distance itself from an interfering object and to yaw freely around an axis without bumping into the interfering object.

48. The method as defined in Claim 45, wherein the first predetermined angle is about 90 degrees.

49. A method of detecting that a mobile robot has been kidnapped, the method comprising:

- receiving an indication that the mobile robot is not instructed to be moving;
- receiving data for video images from a camera coupled to the mobile robot;
- comparing data from different video images to determine whether or not the mobile robot is in motion; and
- determining that the mobile robot has been kidnapped when the video images indicate that the mobile robot is in motion.

50. A circuit for a mobile robot that is configured to determine a motional state of the mobile robot, the circuit comprising:

- a means for receiving pixel data for video images, where the video images are taken from a camera mounted to the mobile robot;
- a means for processing the pixel data for the video images to identify amounts of spatial gradient within a video image;
- a means for characterizing pixels of a video image into at least a first group and a second group, wherein the pixels of the first group correspond to a higher spatial gradient than the pixels of the second group; and
- a means for using the characterization of the pixels to compare a first video image to a second video image to detect the motional state of the mobile robot.

51. A computer program embodied in a tangible medium for controlling a behavior of a mobile robot based on a mismatch between an intended motional state and a perceived motional state, the computer program comprising:

- a module with instructions for receiving an indication of the intended motional state, where the motional state is selected from the group including moving and not moving;

- a module with instructions for using visual data from a camera that is coupled to the mobile robot to perceive the motional state of the mobile robot, where the perceived motional state of the mobile robot is selected from the group including moving and not moving;

- a module with instructions for comparing the intended motional state to the perceived motional state to detect whether a mismatch exists between the intended motional state and the perceived motional state; and

- a module with instructions for changing the behavior of the mobile robot at least partly in response to a detected mismatch.

52. A circuit for control of a self-navigating mobile robot, the circuit comprising:

- a circuit configured to receive an indication that the mobile robot is intended to be traveling in a forward direction;

- a circuit configured to determine from visual image data collected from a video camera coupled to the mobile robot that the mobile robot has ceased traveling in a forward direction;

- a circuit configured to discontinue commands to propel the mobile robot in the forward direction;

- a circuit configured to command the mobile robot to travel in a reverse direction for at least a predetermined distance;

- a circuit configured to determine that the mobile robot has traveled in the reverse direction for at least about the predetermined distance;

- a circuit configured to discontinue commands to propel the mobile robot in the reverse direction;

- a circuit configured to instruct the mobile robot to yaw by at least a first predetermined angle; and

a circuit configured to command the mobile robot to resume forward motion.